Fermilab

Title: Lattice QCD Production on Commodity Clusters at Fermilab

Large scale QCD Monte Carlo calculations have typically been performed on either commercial supercomputers or specially built massively parallel computers. Commodity clusters equipped with high performance networking equipment present an attractive alternative, achieving superior performance to price ratios and offering clear upgrade paths.

We describe the construction and results to date of Fermilab's three Myrinet-networked production clusters (an 80-node dual Pentium III cluster, a 48-node dual Xeon cluster, and a 128-node dual Xeon cluster), as well as an 8-node prototype Itanium 2 cluster. The Xeon and Itanium 2 clusters were funded through the DOE SciDAC "National Infrastructure for Lattice Gauge Computing" project.

Lattice QCD codes are memory bandwidth, floating point, and network intensive. Successful design of clusters to run these codes requires knowledge of the bottlenecks which control performance. Using the MILC lattice QCD code, we examine a number of aspects of performance. For single systems, we discuss the effects of the various memory architectures. We examine optimizations such as data layout and SSE-assisted matrix algebra, and consider SMP behavior. We discuss scaling of the MILC code on Myrinet-connected clusters. Using a modified version of GM which allows control of bandwidth and latency, we examine the sensitivity of the code to network performance.

- Don Holmgren, Paul Mackenzie, Don Petravick, Amitoj Singh, Jim Simone